Security Challenges in Supporting Grid Computing and Collaboration

Mary R. Thompson, Keith R. Jackson
Distributed Systems Department – LBNL

Barriers to adoption

The major challenge for distributed security researchers is to produce solutions that meet the fundamental goals of protecting DOE sites from malicious users and ensuring accountability for use of resources, but still enable sharing of resources, data and collaborative spaces between a broad range of "authorized" users. While DOE security policy may be broadly compatible with shared use of resources and close collaboration among DOE researchers and colleagues, current site security systems implementing those policies often present obstacles to both Grid computing and collaborative technologies. Some examples are:

- Requiring users to have an individual account at each site they access.
- Statically configured firewalls that only open a restricted set of ports or accept connections from previously vetted sites.
- Suspicion or rejection of P2P communication protocols and open source file sharing programs.
- The need for two-factor authentication on all connections to prevent identity theft.

These issues need to be addressed both in the short term to allow the current collaboratories to expand to more sites, and in the long-term to ensure that developing Grid software can interact with future site security systems. Communication and cooperation between those responsible for ensuring site security and those designing and implementing security middleware for distributed applications is essential for the development of middleware that can overcome these obstacles.

The Road Ahead

Ensuring security for distributed computing can be divided into several fairly independent areas:

- Identity management and authentication of users
- Secure communication between authenticated users
- Authorization policy and enforcement of access to resources
- Privacy of user information
- Auditing and accountability of resource use

We will briefly touch on each area to see what the state of the art is and what the possible future directions might be.

Identity and Authentication

Kerberos is an early and still viable solution to global identity and authentication in somewhat restricted environments. PKI identities combined with the TLS protocol is the solution widely adopted in the Grid. Federated identity and attribute servers are the solution used by Internet2 sites [1] and some commercial Web sites. Each of these has strengths and weakness. Sites which already had Kerberos installed have made efforts to integrate Kerberos identities with Grid identities [2]. The PKI solution has spawned a number of efforts to make user credential management more tractable [3]. Federated identity/attribute servers have been developed for individual virtual organizations especially within the EDG [4].

The use of a common format credential with Grid middleware, browser middleware and site authentication is not yet seamless. Rather there are many different solutions to different parts of the problem. One of the reasons for so many solutions is that authentication is the first critical step to any trusted use of resources. Grid authentication must also interact with user and resource site authentication requirements, so a single monolithic solution such as public key credentials or Kerberos is probably not feasible. The idea of federating authentication from various servers and mapping credentials between a common Grid format and site specific credentials is the most promising current solution. This is an example of work that needs to be

both closely tied to real sites and users and to have a broad perspective so that it can produce solutions that will work for many sites and users.

Secure communication

Secure connections have been dominated by TSL type solutions [5]. One of the major reasons for scientists to adopt Globus toolkit solutions to distributed computing was the GSI library that built on SSL and added the delegation required to run parts of jobs at several remote sites. This is an area where some fundamental research is needed as the protocol gets extended to higher levels and needs to accommodate password (and one time password) authentication. The need for secure group communication for the AccessGrid and other group applications also forces the extension of secure communication technology [6].

Authorization

Several collaboration-funded efforts have addressed the authorization requirements of distributed computing and collaboration [7,8,9]. The requirements vary drastically here depending on the application. Often the actual authorization requirements are deceptively simple and all the complications arise in the naming of users and resources and who should be allowed to set the authorization policy. While authentication systems must be simple for the users, but not necessarily the administrators, access policy should be transparent to the users of resources and easy to set and maintain by the owners of individual resources and site administrators. There is also a considerable gap between the authorization requirements of sites providing valuable resources and looser short-lived collaboration between researchers. The former typically require strict individual accountability for auditing and billing of resource use, pre-contracting for resource use, and may provide quality of service guarantees. They can tolerate fairly static policy that is administered by a small number of people. Lightweight collaborative environments, on the other hand, need to be easy to set up and administer, benefit greatly from policy that supports incremental and variable trust of entities and support delegation of rights between users.

This area needs both short –term development to more completely integrate current VOs with existing site requirements and fundamental research into more general solutions. Such research should include: machine parsable policy languages and how to relate them to policy that stakeholders can understand and write; light-weight and portable authorization for ad hoc (such as P2P) and established group collaborations. Another recurrent authorization issue is how to name users in a manner that is meaningful both at the resource site and across the Grid. The adoption of federated identity attribute servers may provide some solutions here.

Privacy

Privacy of user information is an issue that has been addressed within the Internet2 [1] and federated identity services communities, but has not received a lot of attention from DOE research. This is partly because the federating of identity and attribute information from many sources exacerbates the problem, as does participation from individuals who do not entirely trust the community or the service providers. If DOE Grids adopt elements of federated identity and wish to expand the community further into academic and foreign domains, guarantees of privacy may be required. This is still an area for basic research as well as establishing what reasonable and supportable expectations should be.

Accounting

Accounting and auditing are fundamental requirements of large sites. One solution has been to map the Grid user to a site local account which is audited like any other user. This begs the problem of auditing a user's or project's use of resources across the Grid, allowing users to access different sites transparently and reporting back to a home institution a users "bill". Some discussion has taken place at GGF on Grid requirements in this area, but solutions are still in the early stages [10]. In the context of followup and recovery from security incidents there are the cross site, or grid-wide, issues of sharing audit information, lookup of contact information for Grid user id's, etc. This area needs a clear understanding of both site requirements and VO requirements, followed by some fundamental design as to how to pass identifying tags such as Grid job ids and Grid user ids within the middleware components.

References

1. **Shibboleth** http://shibboleth.internet2.edu/
2. **Kerberos CA** "Use of Kerberos-Issued Certificates at FermiLab", M. Crawford & D. Skow, http://chep03.ucsd.edu/files/422.ppt, Presentation at CHEP03
3. **MyProxy** http://grid.ncsa.uiuc.edu/myproxy/
4. **VOMS** "Managing Dynamic User Communities in a Grid of Autonomous Resources" (presented at CHEP 2003 Conference, La Jolla, March 24-28, 2003) http://grid-auth.infn.it/docs/voms_chep2003.pdf
5. **TLS** http://www.openssl.org
6. **SGL** "An Integrated Solution for Secure Group Communication in Wide-Area Networks" (paper and presentation), D. A. Agarwal, O. Chevassut, M. R. Thompson and G. Tsudik *Proceedings of the 6th IEEE Symposium on Computers and Communications*, Hammamet, Tunisia, July 3-5, 2001 http://dsd.lbl.gov/SecGrpComm/Publications/ISCC01/iscc01.html
7. **Akenti** "Certificate-based Authorization Policy in a PKI Environment". M.Thompson, A. Essiari, S. Mudumbai. ACM Transactions on Information and System Security, (TISSEC), Volume 6, Issue 4 (November 2003) http://dsd.lbl.gov/Akenti
8. **CAS** "A Community Authorization Service for Group Collaboration." L. Pearlman, V. Welch, I. Foster, C. Kesselman, S. Tuecke. *Proceedings of the IEEE 3rd International Workshop on Policies for Distributed Systems and Networks*, 2002. http://www.globus.org/research/papers/CAS_2002_Revised.pdf
9. **PRIMA** – Privilege Management and Authorization System, http://zuni.cs.vt.edu/publications/PRIMA-2003.pdf
10. **GGF SAAAR RG** – "Grid Authentication Authorization and Accounting Requirements Draft 5 "– May 21,2004 https://forge.gridforum.org/projects/saaa-rg/document/draft-ggf-saaar-reqs-5.txt/en/1